Computing as Seen by the Experiments: Successes, Failures, Wishes
Introduction to the Special Opening Session

Thomas Nash
Fermi National Accelerator Laboratory
Batavia, IL 60510

This special opening session featured four experimentalists with recent experience producing physics using the computing machinery and tools available to them and their groups. They represented large and successful collaborations working on fixed target and colliding beam experiments in Europe and America. They described their physics results, how they used computing to get them, what went right and what went wrong, and their informed wishes for the future. Their presentations were followed by a provocative audience discussion focussed on the primary goal of the CHEP series: how to get physics done most effectively.

The speakers were Manuel Delfino from Barcelona, a member of the Aleph collaboration at CERN; Joel Butler from Fermilab and the fixed target charm experiment 687; Cathy Newman-Holmes from Fermilab and the CDF collaboration; and Giles Barr, CERN, representing the NA31/NA48 kaon experiments.

The intention in organizing this session was to emphasize the *HEP* in the conference name, CHEP. It was hoped that the conference participants would be encouraged at this session to focus during the rest of the week on evaluating the new computing technologies being disussed in terms of how effectively they enable high energy physics.

Common among the four speakers were calls for better tools, hardware, operating systems, languages, etc. In all of these areas there exist opportunities to significantly improve the ability to do physics, but the transition to new technologies in large collaborations is uniformly seen as a serious problem. The difficulties result from many factors associated with technology injection including a natural conservative resistance to the learning and reprogramming investment associated with new approaches. Often there are extremely complex infrastructures built into large experiments involving both hardware and software that pose formidable barriers to change. One wonders whether transitions can only be made as new collaborations are formed at the beginning of a 10 - 15 year life cycle of a new detector. Are the people who will soon determine the standards and technologies for LHC and SSC detectors establishing the technologies we will use in 2010 AD? Is there some way to design in the ability to evolve technologies through the life of an experiment detector? It seems a terrible shame not to be able to take advantage of the rapid improvements in all areas of computer technology that we can expect over the next 20 years.

Another area of concern noted by several of the speakers has to do with collaborative development of software and publications. Collaborations struggle with preparing the papers that are the primary product of the analysis of their data. It is difficult to find satisfactory standards — and promulgate them throughout a collaboration — that allow for both sophisticated mathematical formula typography and the insertion of graphic objects. Not only do the documents need to be distributed rapidly by network throughout the collaboration, but there is a strong need for support of remote annotations and edits, version management and approval, etc.

Similarly, software development is carried out by hundreds of "amateur" physicists who are geographically dispersed. Here there are additional issues associated with the use of a variety of platforms and compilers, with code certification and testing. Many of the needs expressed with regard to collaborative document and software preparation are being addressed by an emerging new class of collaborative tool support referred to as groupware. When the right groupware tools are identified, they are likely to win rapid acceptance by collaborations — without encountering the kind of transition investment problems referred to earlier.

However, groupware will not address many collaborative software issues such as those associated with the desire to more effectively reuse the work of others (basic elements of reconstruction routines, for example) and problems associated with one physicist's code interfering unintentionally with another's. Many believe that object oriented programming tools are the answer, with inheritance and encapsulation directly attacking these matters. But the question of introducing OOPS into widespread collaboration use brings us back once again to the fundamental and very difficult challenge of how to inject such paradigm shifting technologies into the ongoing work of our detector groups.

The papers from this session are published in the following pages.